

CORRECTING CAMERA SHAKE BY INCREMENTAL SPARSE APPROXIMATION

Paul Shearer, Anna C. Gilbert, Alfred O. Hero III

University of Michigan, Ann Arbor

ABSTRACT

The problem of deblurring an image when the blur kernel is unknown remains challenging after decades of work. Recently there has been rapid progress on correcting irregular blur patterns caused by camera shake, but there is still much room for improvement. We propose a new blind deconvolution method using incremental sparse edge approximation to recover images blurred by camera shake. We estimate the blur kernel first from only the strongest edges in the image, then gradually refine this estimate by allowing for weaker and weaker edges. Our method competes with the benchmark deblurring performance of the state-of-the-art while being significantly faster and easier to generalize.

1. INTRODUCTION

In the problem of blind image deconvolution, we are given a blurry image y and challenged to determine an estimate x of the unknown sharp image x^{true} without knowledge of the blur kernel k^{true} . In the simplest model of blur, y is formed by convolving x^{true} with k^{true} and adding noise n :

$$y = k^{\text{true}} * x^{\text{true}} + n. \quad (1)$$

This convolution model assumes spatially uniform blur, which is frequently violated due to slight camera rotations and out-of-plane effects [1]. Still, the uniform model works surprisingly well and methods for it can be extended to handle nonuniform blur [2, 3].

Even with uniform blur by a single kernel, the blind deconvolution problem is highly underdetermined and additional assumptions must be made to obtain a solution. These assumptions are often imposed most conveniently by moving the problem into a filter space. We define filters $\{f_\gamma\}_{\gamma=1}^L$ and set $y_\gamma = f_\gamma * y$ and $x_\gamma^{\text{true}} = f_\gamma * x^{\text{true}}$, so that

$$y_\gamma = k^{\text{true}} * x_\gamma^{\text{true}} + n_\gamma \quad (2)$$

for $\gamma \in [L] = \{1, \dots, L\}$. Defining $\mathbf{x}^{\text{true}} = \{x_\gamma^{\text{true}}\}_{\gamma=1}^L$, $\mathbf{y} = \{y_\gamma\}_{\gamma=1}^L$, and $(k * \mathbf{x}^{\text{true}})_\gamma = k * x_\gamma^{\text{true}}$, we can write the filter space problem compactly as

$$\mathbf{y} = k^{\text{true}} * \mathbf{x}^{\text{true}} + \mathbf{n}. \quad (3)$$

The simplest nontrivial filter space is gradient space, where $L = 2$ and $f_1 = [1, -1]$, $f_2 = [1, -1]^T$, but there are many

other possibilities. Determining x from a filter space representation \mathbf{x} often does not work well, so typically one obtains an estimate k of k^{true} and deconvolves y with k to get x [1].

Bayesian inference is a convenient framework for imposing prior assumptions to regularize blind deconvolution [4]. By assuming some distribution of \mathbf{n} we obtain a likelihood function $p(\mathbf{y} | k * \mathbf{x})$ which gives the probability that the data \mathbf{y} arose from a given pair (k, \mathbf{x}) . We then choose priors $p(k)$ and $p(\mathbf{x})$ and compute the posterior distribution

$$p(k, \mathbf{x} | \mathbf{y}) \propto p(\mathbf{y} | k * \mathbf{x})p(\mathbf{x})p(k). \quad (4)$$

Estimates of \mathbf{x} and k may be obtained by summary statistics on $p(k, \mathbf{x} | \mathbf{y})$. We call the mode of $p(k, \mathbf{x} | \mathbf{y})$ the joint maximum *a posteriori* (MAP) estimator, while the mode of the marginal $p(k | \mathbf{y}) = \int p(k, \mathbf{x} | \mathbf{y}) d\mathbf{x}$ is the kernel MAP estimator. Most blind deconvolution methods are nominally MAP estimators but do not actually find a global minimizer, as this is typically intractable and may even be counterproductive. We refer to any method organized around optimizing a posterior as a MAP method, while methods that actually find a global minimum will be called *ideal* MAP methods. Joint MAP methods typically attempt to minimize the cost function $F(k, \mathbf{x}) = -\log p(k, \mathbf{x} | \mathbf{y})$, which may be written (up to an irrelevant additive constant) as the sum of a data misfit and two regularization terms,

$$F(k, \mathbf{x}) = L(k * \mathbf{x}) + R_{\mathbf{x}}(\mathbf{x}) + R_k(k), \quad (5)$$

where each of these functions may take the value $+\infty$ to represent a hard constraint. Kernel MAP estimation is more difficult as it involves a high-dimensional marginalization, and it is typically approximated by variational Bayes or MCMC sampling [5].

Joint MAP estimation is the oldest, simplest, and most versatile approach to blind deconvolution [6–8], but initial joint MAP efforts on the camera shake problem met with failure [9], even when ℓ_p regularizers for $p < 1$ were used. In [1], Levin *et al* showed that the ℓ_p regularizer generally prefers blurry images to sharp ones: $\|\mathbf{y}\|_p^p < \|\mathbf{x}^{\text{true}}\|_p^p$, so that ideal joint MAP typically gives the trivial *no-blur* solution $(k, x) = (\delta_0, y)$, where δ_0 is the Kronecker delta kernel. The non-ideal joint MAP methods [10, 11] somewhat compensate for the defects of ideal joint MAP by dynamic edge prediction and likelihood weighting, but benchmarking in [1, 12] showed that these heuristics sometimes fail.

In [9] Fergus *et al* developed a kernel MAP method with a sparse edge prior which was very effective for correcting camera shake. In [1] it was noted that marginalization over \mathbf{x} seems to immunize ideal kernel MAP against the blur-favoring prior problem. More refined kernel MAP methods were recently reported in [12] and [13], and to our knowledge these two methods are the top performers on the benchmark 32 image test set from [1]. While these efforts have made kernel MAP much more tractable, it remains harder to understand and generalize than joint MAP, so it would be useful to find a joint MAP method that is competitive with kernel MAP on the camera shake problem.

In [14], Krishnan *et al* addressed the blur-favoring prior problem in joint MAP by changing the prior, proposing the scale-invariant ℓ_1/ℓ_2 ratio as a ‘normalized’ sparse edge penalty. The ℓ_2 normalization compensates for the way that blur reduces total ℓ_1 edge mass, causing the ℓ_1/ℓ_2 penalty to prefer sharp images and eliminating the need for additional heuristics. While their algorithm does not quite match the performance of [9] on the benchmark test set from [1], it comes fairly close while being significantly simpler, faster, and in some cases more robust. Other promising joint MAP methods include [15–17], but we are not aware of public code with full benchmark results for these methods.

1.1. Our approach

We propose a new approach to joint MAP blind deconvolution in which the kernel is estimated from a sparse approximation \mathbf{x} of the sharp gradient map \mathbf{x}^{true} . Initially we constrain \mathbf{x} to be very sparse, so it contains only the few strongest edges in the image, and we determine k such that $k * \mathbf{x} \approx \mathbf{y}$. Because \mathbf{x} is so much sparser than \mathbf{y} , the solution $k = \delta_0$ is very unlikely. But generally this initial k overestimates k^{true} , so we refine k by letting weaker edges into \mathbf{x} .

To present this approach formally, we set $f_1 = [1, -1]$, $f_2 = [1, -1]^T$, so that $\mathbf{x}(p) = (x_1(p), x_2(p))$ is the discrete image gradient vector at each pixel p . We set $L(k, \mathbf{x}) = \frac{1}{2} \|k * \mathbf{x} - \mathbf{y}\|_2^2$ and impose the usual positivity and unit sum constraints on k . We measure gradient sparsity using the $\ell_{2,0}$ norm: $|\mathbf{x}(p)|$ is the ℓ_2 length of $\mathbf{x}(p)$ and $\|\mathbf{x}\|_{2,0} = \|\mathbf{x}\|_0$ the number of nonzero gradient vectors. The joint MAP optimization problem is then

$$\begin{aligned} & \underset{k, \mathbf{x}}{\text{minimize}} && \frac{1}{2} \|k * \mathbf{x} - \mathbf{y}\|_2^2 \\ & \text{subject to} && k \geq 0, \mathbf{1}^T k = 1, \|\mathbf{x}\|_{2,0} \leq \tau, \end{aligned} \quad (6)$$

where the expression $a^T b$ denotes the dot product of the arrays a and b when considered as vectors, and the $\mathbf{1}$ in $\mathbf{1}^T k$ is an all-ones array.

We solve this problem with an iterative optimizer described in §2, and slowly increase τ as the iterations proceed. To initialize τ we use the ℓ_1/ℓ_2 ratio, a robust lower bound

on a signal’s sparsity [18]. The sharp \mathbf{x} should be significantly sparser than \mathbf{y} , so initially we set $\tau = \beta_0 \tau_{\mathbf{y}}$, where $\tau_{\mathbf{y}} = \|\mathbf{y}\|_1 / \|\mathbf{y}\|_2$ and $\beta_0 < 1$ is a small constant. After an initial burn-in period of I_b iterations we multiply τ by a constant growth factor $\gamma > 1$, an action we repeat every I_s iterations thereafter.

We use a standard multiscale seeding technique to accelerate the kernel estimation step [9, 14]. We begin by solving (6) with a heavily downsampled \mathbf{y} , giving a cheap, low-resolution approximation to k and \mathbf{x} . We then upsample this approximation and use it as an initial guess to solve (6) with a higher resolution \mathbf{y} , repeating the upsample-and-seed cycle until we reach the full resolution \mathbf{y} . At each scale we use the same τ increase schedule. After kernel estimation we use non-blind deconvolution of \mathbf{y} with k to get the sharp image \mathbf{x} .

The easiest way to understand how our kernel estimation works is to watch k and \mathbf{x} evolve as the iterations progress. In Fig. 1, the state of k and \mathbf{x} is shown at iterations 2, 32, and 150 of the final full-resolution scale, with k^{true} and \mathbf{x}^{true} at far right. Initially \mathbf{x} is quite sparse, so k cannot be a trivial kernel because the parts of \mathbf{y} not in \mathbf{x} must be attributed to blur. But this initial approximation is crude, so as τ increases with iteration, \mathbf{x} is allowed to have more and more edges so that k can be refined.

1.2. Novelty and relations with existing methods

Direct ℓ_0 optimization is well-established in the compressed sensing community [19, 20] but we are not aware of any effective ℓ_0 approaches to blind deconvolution. In [14] the ℓ_1/ℓ_2 ratio was deliberately chosen over ℓ_0 because while both have the desired scale invariance, the graph of ℓ_1/ℓ_2 is smoother and looks more ‘optimizable’ than ℓ_0 . We contend that ℓ_0 may be difficult to use as a cost function, but very effective as a constraint. Gradient and kernel thresholding are commonly used [10, 11] and these can be interpreted as ℓ_0 projections, but they are typically used as auxiliary heuristics, not as the central modeling idea. Our technique of slowly increasing the sparsity constraint τ is reminiscent of matching pursuit algorithms for sparse approximation [21, 22]. It is also related to the likelihood reweighting technique of [10], which may be seen by considering the Lagrangian of (6). However, our initialization strategy requires that we use the constrained formulation rather than the Lagrangian.

2. ALTERNATING PROJECTED GRADIENT METHOD

To solve problem (6) at a given scale, we use a standard alternating descent strategy: starting from some initial k and \mathbf{x} , we reduce the cost function by updating \mathbf{x} with k fixed, then k with \mathbf{x} fixed, cycling until a stopping criterion is met. Each cycle, or outer iteration, consists of $I_{\mathbf{x}}$ inner iterations updating \mathbf{x} and I_k inner iterations updating k . All updates are



Fig. 1. Kernel estimation on an image from the test set of [1]; a small patch has been selected and rescaled for clarity. *Left:* Blurry edge map $|\mathbf{y}|$. *Center left to center right:* Evolution of the kernel k (inset) and edge map magnitude $|\mathbf{x}|$ in the final full-resolution stage. As τ increases in (6), the edge map becomes less sparse and the kernel is refined. *Right:* k^{true} and $|\mathbf{x}^{\text{true}}|$.

computed with a projected gradient method; given a smooth function $h(u)$ and a constraint set \mathcal{U} , projected gradient methods seek a solution of $\min_{u \in \mathcal{U}} h(u)$ by updates of the form $u \leftarrow \mathcal{P}_{\mathcal{U}}(u - \alpha_u g_u)$, where $g_u = \nabla h(u)$, α_u is a step size, and $\mathcal{P}_{\mathcal{U}}(w) = \operatorname{argmin}_{u \in \mathcal{U}} \|u - w\|_2^2$ is the Euclidean projection of w onto \mathcal{U} . Convergence of alternating descent and projected gradient methods to stationary points is proven in [23] under mild conditions.

We now describe how we compute the projected gradient iterations for the inner subproblems $\min_{k \in \mathcal{K}} L(k, \mathbf{x})$ and $\min_{\mathbf{x} \in \mathcal{X}} L(k, \mathbf{x})$, where $L(k, \mathbf{x}) = \frac{1}{2} \|k * \mathbf{x} - \mathbf{y}\|_2^2$, $\mathcal{K} = \{k \mid k \geq 0, 1^T k = 1\}$, and $\mathcal{X} = \{\mathbf{x} \mid \|\mathbf{x}\|_{2,0} \leq \tau\}$. Letting $\mathbf{r} = k * \mathbf{x} - \mathbf{y}$ denote the residual, we have $\nabla_k L = \sum_{\gamma} \bar{\mathbf{x}}_{\gamma} * \mathbf{r}_{\gamma}$ and $\nabla_{\mathbf{x}} L = \bar{k} * \mathbf{r}$, where the bar denotes 180° rotation about the origin. Assuming the nonzero elements of $|\mathbf{x}|$ are distinct, the projection $\mathcal{P}_{\mathcal{X}}(\mathbf{x})$ is the top- τ vector thresholding

$$\mathcal{P}_{\mathcal{X}}(\mathbf{x})(i) = \mathbf{x}(i) \cdot \mathbf{1}(|\mathbf{x}(i)| \geq \theta(|\mathbf{x}|, \tau)), \quad (7)$$

where $\mathbf{1}(A)$ is the indicator function for condition A and $\theta(|\mathbf{x}|, \tau)$ is the τ^{th} biggest element of $|\mathbf{x}|$. The set \mathcal{K} is a canonical simplex with projection $\mathcal{P}_{\mathcal{K}}(k)$ given by

$$\mathcal{P}_{\mathcal{K}}(k)(i) = \max(0, k(i) - \sigma), \quad (8)$$

where σ is the unique solution of $1^T \mathcal{P}_{\mathcal{K}}(k) = 1$. Both $\mathcal{P}_{\mathcal{X}}$ and $\mathcal{P}_{\mathcal{K}}$ can be computed in linear time using selection algorithms [24, 25].

The step sizes $\alpha_{\mathbf{x}}, \alpha_k$ are chosen by backtracking line search from an initial guess. In the \mathbf{x} subproblem our initial guess is

$$\alpha_{\mathbf{x}} = \frac{(k * g_{\mathbf{x}})^T \mathbf{r}}{(k * g_{\mathbf{x}})^T (k * g_{\mathbf{x}})}, \quad (9)$$

which is optimal in the sense that it solves the problem $\min_{\alpha} L(k, \mathbf{x} - \alpha g_{\mathbf{x}})$. This aggressive step size was chosen over several alternatives, as it was the most effective for securing good edge support estimates. In the k subproblem we use the spectral projected gradient (SPG) method [26]; in the first iteration $\alpha_k = 1$, and in subsequent iterations we use the

Barzilai-Borwein step size

$$\alpha_k = \frac{(g_k - g_k^{\text{old}})^T (g_k - g_k^{\text{old}})}{(g_k - g_k^{\text{old}})^T (k - k^{\text{old}})} \quad (10)$$

where g_k^{old} and k^{old} denote the values of g_k and k at the previous SPG iteration.

3. IMPLEMENTATION AND EXPERIMENTS

We implemented our method in MATLAB by modifying the code of [14], which uses a similar strategy of alternating minimization with multiscale seeding. The full-resolution kernel size was set to 35×35 for all experiments. The initial stage of the multiscale algorithm downsamples \mathbf{y} by a factor of $5/35$ in each direction, so that the kernel is of size 5×5 , and each upsample cycle increases the size of k , \mathbf{x} , and \mathbf{y} by a factor of roughly $\sqrt{2}$ until full resolution is reached. The parameters of the core single-scale algorithm from §2 were set to $\beta_0 = 0.15$, $\gamma = 1.10$, $I_b = 20$, $I_s = 10$, $I_{\mathbf{x}} = 1$, $I_k = 6$. We do 30 iterations of the alternating projected gradient method for all scales except the final, full-resolution scale, which uses 180 iterations. Non-blind deconvolution with the estimated kernel was performed using the method of [27], using the parameter settings chosen in the code for [12].

In [1] a test set of 32 blurry images with known ground truth was created for benchmarking blind deconvolution methods. Each blurry image was formed by taking a picture of a sharp image with a camera that shook in-plane, and bright points outside the image were used to obtain ground truth blur kernels. A total of 32 blurry images were formed by blurring 4 sharp images on 8 different shake trajectories. This test set has become the *de facto* standard for objectively comparing different methods.

We ran our algorithm on this test set and compared its performance against the methods of [12] and [13]. We compare against these methods because they have published implementations which match or exceed the performance of the state-of-the-art methods in [9–11, 14], and we know of no




Fig. 2. Sample results from our method, [12, 13] on the benchmark set of [1]. True and recovered kernels inset.




Fig. 3. Cumulative deblurring performance our method, [12], and [13] on the 32 image test set of [1]. The vertical axis is the percentage of the 32 runs having at most a given SSE.

methods that outperform [12] and [13] on this test. We use the squared error metric $SSE(x) = \sum_i (x(i) - x^{\text{true}}(i))^2$ to measure performance and note that results using the ratio metric of [1] are similar. Results for [12] were taken from files included with their published implementation, while results for [13] were generated by running their online test script using the log prior, which was the best in their experiments.

Our experiments were performed in MATLAB 2011b on an Intel Quad Core Xeon 2.2 GHz Mac Pro. Our method’s kernel estimation step took 45 – 60 seconds per image, and deconvolution took 15 seconds. The other methods took 45 – 240 seconds for kernel estimation, and their computation time depended strongly on kernel size. The difference is mostly due to our use of cheap SPG iteration rather than quadratic programming in the k step, and also because \mathcal{P}_K and \mathcal{P}_X make k and x sparse, enabling $k * x$ to be computed faster.

Sample results on the benchmark are shown in Fig. 2. Our method and [12] perform very similarly on both images. On the house image [13] is very sharp and by far the best, but it suffers from severe artifacts on the boy image. Fig. 3 summarizes the full-benchmark performance of the three methods

using cumulative error histograms. The curves for our method and [12] are largely similar. The curve for [13] is above ours and [12] for about half of the images, but it flattens out below 85% while the others plateau at 100%. This is because method [13] struggled on the boy images. We note that the results reported in [13] for this test set are better than those obtained in our run of their code, although we ran it without any modification. The authors of [13] state that their code is a simplification of what was used to generate the reported results. While there may be a more sophisticated version of their code that outperforms ours, our method competes with the available state of the art.

4. CONCLUSION

We have proposed a blind deconvolution method in which the blur kernel is estimated by incremental sparse edge approximation. A rough global blur kernel is first estimated from only the strongest edges in the image, then it is refined as we allow the image edge map to gradually become less and less sparse. Ours is the first simple, fast joint MAP method to match the state-of-the-art kernel MAP methods in [12, 13] on an objective benchmark. The success of the methods in [14] and this paper suggest that the downsides of ideal joint MAP described in [1] can be robustly avoided without resort to a more complex kernel MAP estimation.

There are many potential avenues for improving and extending our method. The edge sparsity relaxation schedule we use is slow and conservative, and a more adaptive schedule could make the method faster. Our initialization of the edge map sparsity does not take noise into account, and may need to be modified for very noisy images. Extension to nonuniform blur models, nonquadratic likelihoods, and fast parallel or GPU implementations are possible. The speed of our kernel estimation may make it useful as an input to high-quality non-blind methods such as [13].

5. REFERENCES

- [1] A. Levin, Y. Weiss, F. Durand, and W.T. Freeman, "Understanding and evaluating blind deconvolution algorithms," *IEEE Conference on Computer Vision and Pattern Recognition*, vol. 0, pp. 1964–1971, 2009.
- [2] S. Cho, Y. Matsushita, and S. Lee, "Removing non-uniform motion blur from images," in *IEEE 11th International Conference on Computer Vision*, 2007, pp. 1–8.
- [3] O. Whyte, J. Sivic, A. Zisserman, and J. Ponce, "Non-uniform deblurring for shaken images," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2010, pp. 491–498.
- [4] P. Campisi and K. Egiazarian, *Blind Image Deconvolution: Theory and Applications*, CRC Press, 2007.
- [5] C.M. Bishop et al., *Pattern Recognition and Machine Learning*, Springer-Verlag New York, 2006.
- [6] G.R. Ayers and J.C. Dainty, "Iterative blind deconvolution method and its applications," *Optics Letters*, vol. 13, no. 7, pp. 547–549, 1988.
- [7] Y. L. You and M. Kaveh, "A regularization approach to joint blur identification and image restoration," *IEEE Transactions on Image Processing*, vol. 5, no. 3, pp. 416–28, 1996.
- [8] T. F. Chan and C.-K. Wong, "Total variation blind deconvolution," *IEEE Transactions on Image Processing*, vol. 7, pp. 370–375, Mar. 1998.
- [9] R. Fergus, B. Singh, A. Hertzmann, S.T. Roweis, and W.T. Freeman, "Removing camera shake from a single photograph," in *ACM Transactions on Graphics*, 2006, vol. 25, pp. 787–794.
- [10] Q. Shan, J. Jia, and A. Agarwala, "High-quality motion deblurring from a single image," in *ACM Transactions on Graphics*, 2008, vol. 27, p. 73.
- [11] S. Cho and S. Lee, "Fast motion deblurring," in *ACM Transactions on Graphics*, 2009, vol. 28, p. 145.
- [12] A. Levin, Y. Weiss, F. Durand, and W.T. Freeman, "Efficient marginal likelihood optimization in blind deconvolution," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2011, pp. 2657–2664.
- [13] S. D. Babacan, R. Molina, M. N. Do, and A. K. Katsaggelos, "Bayesian blind deconvolution with general sparse image priors," in *European Conference on Computer Vision (ECCV)*, Firenze, Italy, October 2012, Springer.
- [14] D. Krishnan, T. Tay, and R. Fergus, "Blind deconvolution using a normalized sparsity measure," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2011, pp. 233–240.
- [15] J.F. Cai, H. Ji, C. Liu, and Z. Shen, "Framelet-based blind motion deblurring from a single image," *IEEE Transactions on Image Processing*, vol. 21, no. 2, pp. 562–572, 2012.
- [16] C. Wang, L.F. Sun, Z.Y. Chen, J.W. Zhang, and S.Q. Yang, "Multi-scale blind motion deblurring using local minimum," *Inverse Problems*, vol. 26, no. 1, pp. 015003, 2009.
- [17] L. Xu and J. Jia, "Two-phase kernel estimation for robust motion deblurring," *European Conference on Computer Vision*, pp. 157–170, 2010.
- [18] M.E. Lopes, "Estimating unknown sparsity in compressed sensing," *arXiv preprint arXiv:1204.4227*, 2012.
- [19] T. Blumensath and M.E. Davies, "Normalized iterative hard thresholding: Guaranteed stability and performance," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 298–309, 2010.
- [20] R. Garg and R. Khandekar, "Gradient descent with sparsification: an iterative algorithm for sparse recovery with restricted isometry property," in *Proceedings of the 26th Annual International Conference on Machine Learning*. ACM, 2009, pp. 337–344.
- [21] J.A. Tropp and A.C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Transactions on Information Theory*, vol. 53, no. 12, pp. 4655–4666, 2007.
- [22] D. Needell and J.A. Tropp, "Cosamp: Iterative signal recovery from incomplete and inaccurate samples," *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301–321, 2009.
- [23] H. Attouch, J. Bolte, and B. Svaiter, "Convergence of descent methods for semi-algebraic and tame problems: proximal algorithms, forward-backward splitting, and regularized gauss-seidel methods," *Mathematical Programming*, pp. 1–39, 2011.
- [24] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, The MIT Press, 3rd edition, 2009.
- [25] J. Duchi, S. Shalev-Shwartz, Y. Singer, and T. Chandra, "Efficient projections onto the ℓ_1 -ball for learning in high dimensions," in *Proceedings of the 25th International Conference on Machine Learning*, 2008, pp. 272–279.
- [26] E.G. Birgin, J.M. Martínez, and M. Raydan, "Nonmonotone spectral projected gradient methods on convex sets," *SIAM Journal on Optimization*, vol. 10, no. 4, pp. 1196–1211, 2000.
- [27] A. Levin, R. Fergus, F. Durand, and W.T. Freeman, "Deconvolution using natural image priors," *Massachusetts Institute of Technology, Computer Science and Artificial Intelligence Laboratory*, 2007.